


# PDSF Users Meeting

- \* PDSF performance
- \* past/future outages
- \* announcements
- \* toward common PDSF Shifter image
- \* PDSF interactive nodes
- \* AOB

July 12, 2016

Jan Balewski

# Utilization



# Past/future outages

## Past

NIM - unavailable last week due to critical bug

## Future

### HPSS:

07/13/16 9:00-13:00 PDT, Scheduled maintenance.

# Announcements

## Bi-weekly office hours in June

- Thursday, July 21, 2-4pm, 59-4016-CR

## PDSF users meeting

- Tuesday, August 2, 11:15 - 12:15pm 59-3034-CR



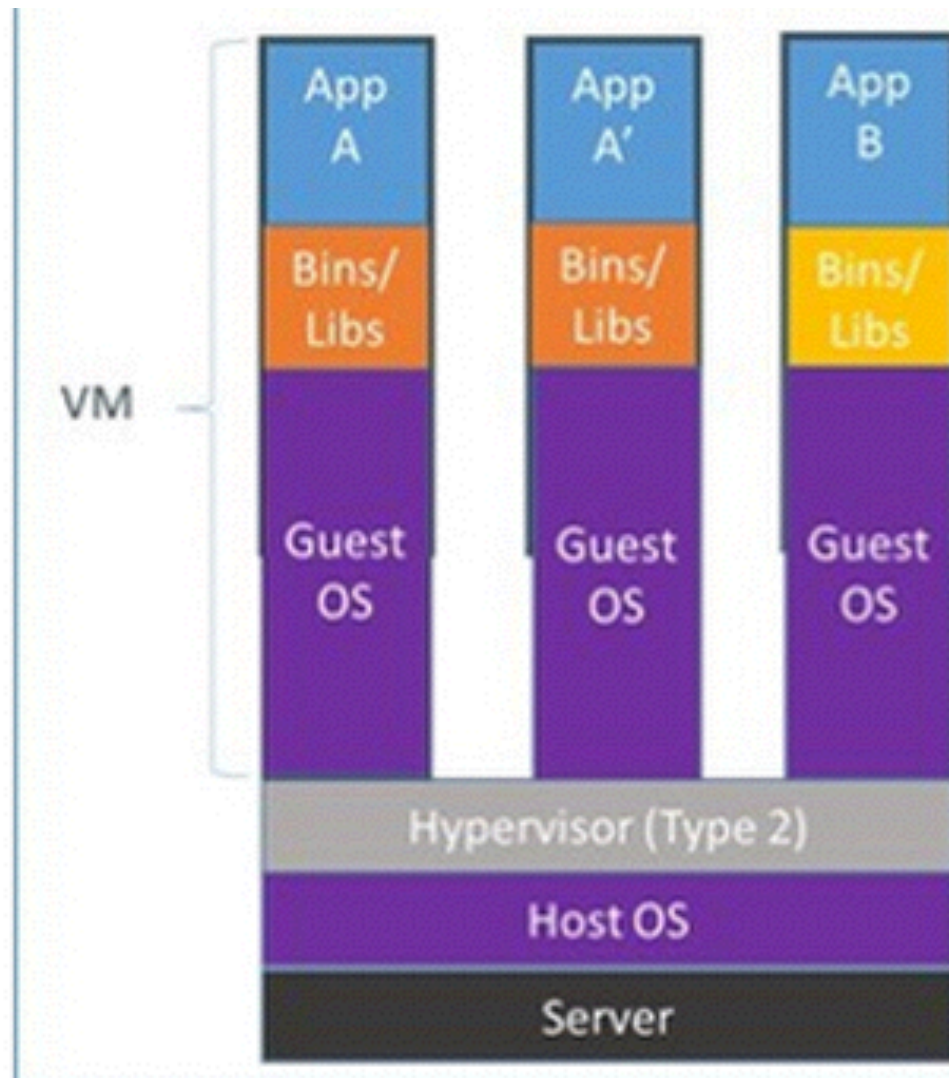
# Toward PDSF-wide Shifter image



August 7, 2015

**NERSC's 'Shifter' Makes Container-based HPC a Breeze**

Kathy Kincade, Communications Specialist, Berkeley Lab Computing Sciences






# Toward PDSF-wide Shifter image

## Research Groups - what do you really need?

**ALICE - A Large**



**STAR - Solenoidal T**



STAR  
of s  
at  
PD  
The  
and

**ATLAS - A Toroid**



**DAYABAY - The D**



**Majorana - Neutrino**



T  
U  
F  
V  
t  
PD  
The

**LUX - Dark Matter S**



pdsf9 \$ module avail

```
pdsf9 $ module avail

----- /usr/share/modules/3.2.10/Modules/versions -----
3.2.10

----- /usr/share/modules/3.2.10/Modules/3.2.10/modulefiles -----
dot          module-git  module-info modules      null          use.own

----- /usr/common/usg/Modules/modulefiles -----
bison/2.7.1(default)      ImageMagick/6.8.9-0(default)  pandas/0.13.1(default)
boost/1.54.0(default)    ipython/1.2.1(default)        python/2.7.10
cernlib/2005(default)    ipython/2.0.0                  python/2.7.3
CLHEP/2.1.0.1             java/jdk1.6.0_10              python/2.7.6(default)
CLHEP/2.1.3.1(default)   java/jdk1.7.0_60(default)     python/2.7.9
CLHEP/2.2.0.8             latex/2013(default)           python/2.7-anaconda
cmake/2.8.12              libunwind/1.1(default)        python/3.4.3
cmake/3.5.1               libxml2/2.9.2(default)        python_base/2.7.3
cmake/3.5.2(default)     libz/1.2.3-32bit-static(default) python_base/2.7.6(default)
curl/7.35.0(default)     m4/1.4.16(default)            qcache grind/0.7.4(default)
emacs/24.5(default)      matlab/R2014a(default)        qt/4.8.6(default)
erlang-OTP/17.1(default) matplotlib/1.2.0               qt/5.5.0
firefox/33.1(default)    matplotlib/1.3.1(default)     ROOT/5.34.10(default)
fluka/2011                mongod/3.0.4                  ROOT/5.34.20
garfield/9(default)      myquota/1.0                   rrdtool/1.4.8(default)
gcc/4.8.1(default)       mysql/5.6.12                   scipy/0.13.3(default)
gcc/4.8.2                mysql++/3.2.1(default)         SGE/8.1.2(default)
Geant4/10.02.b01          mysqlpython/1.2.4(default)    subversion/1.6.11(default)
Geant4/4.9.4.p04          nedit/5.6(default)            subversion/1.8.11
Geant4/4.9.5(default)    numpy/1.7.0                    texLive/2015
git/2.2.1(default)        numpy/1.8.1(default)          tmux/1.9a(default)
globus/3.1.35-tarball-1(default) openmpi/1.6.5(default)        usg/1.0
gperftools/2.2.1         oprofile/0.9.9                usg-default-modules/1.0(default)
gperftools/2.3(default)  osg/3.1                        valgrind/3.9.0(default)
hdf5/1.8.13(default)     osg/3.1.35-tarball-1(default) xpdf/3.03(default)

----- /usr/syscom/nsg/modulefiles -----
gpfs/1.1.0(default)  nsg/1.2.0(default)

----- /usr/syscom/nsg/opt/modulefiles -----
BerkeleyDB/5.1.19(default) cfengine/3.2.3(default)
```


# Toward PDSF-wide Shifter image

## Any volunteers?

### Proof-of-principle STAR Shifter image from 2013

<https://www.nersc.gov/users/computational-systems/pdsf/group-pages/star/shifter-image-of-star-sl13a/>

STAR - Solenoid



```
my-laptop$ ssh -A -X edison.nersc.gov
```

```
edison06:~>$ module load shifter /* load needed module */
```

```
/* request one interactive node with 32 cores for 30 minutes */
```

```
balewski@edison06:~> salloc -N 1 -p debug --image=docker:balewski/sl64:STAR-sl13a4 -t 00:30:00
```

```
/* in order to run STAR simulation/analysis you need to setup manually STAR environment, execute */
```

```
[balewski@nid00009 ~]$ source /usr/local/star/group/templates/cshrc
```

```
----- STAR Group Login from /usr/local/star/group/ -----
```

```
Setting up STAR_ROOT = /usr/local/star/
```

```
Setting up STAR_PATH = /usr/local/star/packages
```

```
Setting up WWW_HOME = http://www.star.bnl.gov/
```

```
...
```

```
/* finally set (the only allowed) STAR library version */
```

```
[nid00009] ~/> starver SL13a
```

```
/* now you can run starsim */
```


```
[nid00009] ~/> starsim
```

```
starsim > detp geom y2009
```

```
starsim > make geometry
```

```
starsim > dcut cave z 1 10 10
```

```
====> Datacard assign R.SENSITIVE ( 1) = 0.0000 in CCUT ba
====> Datacard assign I.ONOFF ( 1) = 1 in EMCG ba
====> Datacard assign I.FILLMODE ( 1) = 3 in EMCG ba
ZCAL
====> Datacard assign R.VERSION ( 1) = 1.0000 in MAGG ba
StarMagField instantiate starsim field=5
StarMagField::ReadField Reading Magnetic Field Measured Fu
StarMagField::ReadField Filename is bfield_full_positive_2D.
StarMagField::ReadField Reading 2D Magnetic Field file: bfi
StarMagField::ReadField Reading 3D Magnetic Field file: bfi
StarMagField::ReadField Reading 3D Magnetic Field file: bfi
StarMagField::SetLock lock StarMagField parameters
StarMagField parameters are locked
StarMagField:: Map 2
StarMagField:: Factor 1.000000
StarMagField:: Rescale 1.000000
StarMagField:: BDipole -42.669998
StarMagField:: RmaxDip 15.340000
StarMagField:: ZminDip 980.000000
StarMagField:: ZmaxDip 1350.000000
StarMagField::mflgeo(5) Bz=4.9798
MFLD magField, Bz = 5.000000 4.9797993
starsim > dcut cave z 1 10 10
starsim >
```



## PDSF interactive nodes

PDSF has 6 interactive nodes: pdsf6,...,11  
each: 32 cores, 128GB RAM, 2.6 GHz clock  
total : 192 cores

pds 9,10,11 were repurposed compute nodes when Mendel was split between OSF and CRT

### Proposal :

- put pdsf 9,10,11 back to compute pool
- use only 3 interactive nodes, 96 cores total



# Proposed limit on use of PDSF interactive nodes

In last 2 months 1 user was reported twice running 5 interactive jobs on the same node for hour+

‘stay below 2’ rule , base don ‘top’ report

- up to 2 jobs with 10%+ CPU (here 5)
- sum of used CPU load <2 (here 5)
- sum of used CPU time < 2 hours (here 21 h)

## pdsf10 load example last week

```
top - 11:54:31 up 47 days,  8:56, 41 users,  load average: 8.51, 8.87, 8.88
Tasks: 1003 total,   6 running, 989 sleeping,   8 stopped,   0 zombie
Cpu(s): 15.8%us,  0.2%sy,  0.0%ni, 84.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Mem:  131665308k total, 42902380k used, 88762928k free,   500852k buffers
Swap: 407199732k total,   609224k used, 406590508k free, 10856432k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
12598	hanseul	20	0	3819m	2.5g	2664	R	100.0	2.0	257:49.56	root.exe
12543	hanseul	20	0	3846m	2.6g	2660	R	99.9	2.0	259:23.42	root.exe
12574	hanseul	20	0	3851m	2.6g	2660	R	99.9	2.0	258:08.51	root.exe
12586	hanseul	20	0	3852m	2.6g	2676	R	99.9	2.0	257:59.44	root.exe
12623	hanseul	20	0	3638m	2.4g	12m	R	99.9	1.9	257:41.91	root.exe

## Proposal: implementation/ violations

- user feedback + consult files ticket
- automatic script :
  - warns (ab)user
  - warns consult
- gradation of consequences
  - 1st warning
  - 2nd warning
  - disable PDSF access in NIM, requires mail form PI to reactivate it

# AOB